האוניברסיטה העברית בירושלים
THE HEBREW UNIVERSITY OF JERUSALEM

THE ROLE OF IMPULSES IN SHAPING DECISIONS

By

JUDITH AVRAHAMI and YAAKOV KAREEV

Discussion Paper # 552 May 2010

מרכז לחקר הרציונליות

CENTER FOR THE STUDY OF RATIONALITY

Feldman Building, Givat-Ram, 91904 Jerusalem, Israel
PHONE: [972]-2-6584135 FAX: [972]-2-6513681
E-MAIL: ratio@math.huji.ac.il
URL: http://www.ratio.huji.ac.il/

The Role of Impulses in Shaping Decisions

Judith Avrahami and Yaakov Kareev[1]

May 10, 2010

ABSTRACT

This article explores the extent to which decision behavior is shaped by short-lived reactions to the outcome of the most recent decision. We inspected repeated decision-making behavior in two versions of each of two decision-making tasks, an individual task and a strategic one. By regressing behavior onto the outcomes of recent decisions, we found that the upcoming decision was well predicted by the most recent outcome alone, with the tendency to repeat a previous action being affected both by its actual outcome and by the outcomes of actions not taken. Because the goodness of predictions based on the most recent outcome did not diminish as participants gained experience with the task, we conclude that repeated decisions are continuously affected by impulsive reactions.

INTRODUCTION

Decisions are often the result of contemplating the outcomes likely to follow them, that is, of calculating the costs and benefits associated with different options; decisions can also be the result of learning from experience the relative merits of alternative options. Either way, decisions are made to improve overall results. However, at times, even choice of the better option may result in an undesirable outcome, and choice of the lesser option in a desirable one. Do such outcomes produce an instinctive reaction that is reflected in the immediately following decision? Furthermore, does the outcome of the most recent decision continue to have an effect even as the expected value of each option becomes increasingly clear? If it does, if the effect is short-lived and its magnitude does not diminish with experience, we shall argue that this short-lived, persistent, reaction is impulsive, the result of an involuntary inclination expressed in the immediately following decision.

Many hold that decisions are shaped by past experience. Opinions differ, however, as to what should be regarded as the *past*, what should be regarded as *experience*, and how past

[1] School of Education and The Center for the Study of Rationality, The Hebrew University of Jerusalem, Jerusalem, Israel

experience *shapes* decisions. Regarding the past, most recent models of learning in decision making (e.g., Busemeyer & Stout, 2002; Camerer & Ho, 1999; Erev & Roth, 1998; Hart, 2005; Roth & Erev, 1995; Selten, Abbink, & Cox, 2005) take the whole set of past decisions (usually with some memory decay) or a small sample drawn from all past experiences (e.g., Erev & Barron, 2005; Selten & Chmura, 2008) into account. Other models, such as the win-stay, lose-shift model (e.g., Levine, 1966; Matsen & Nowak, 2004; Milinski, 1993; Nowak & Sigmund 1993) and the Cournot dynamics model (Cournot, 1838/1960), consider only the very last experience.

Regarding experience, models of decision making all consider the actions taken but differ in what role they assign to the actions not taken. Some, following in the tradition of classical learning theory in psychology (e.g., win-stay, lose-shift model—Levine, 1966; reinforcement learning—Roth & Erev, 1995), consider only the actions actually taken. In these models, the outcome of actions that could have been, but were not, taken are ignored—even if known. In other models, if the outcomes of actions not taken are known, they too are considered in predicting behavior. In such models, it is the comparison between the outcome of an action and the counterfactual outcomes that determines the subsequent attraction of the option taken (Camerer & Ho, 1998, 1999; Hart & Mas-Colell, 2000; Ho, Camerer, & Chong, 2007; Selten et al., 2005; Selten & Chmura, 2008). The latter models differ among themselves in the way they use the comparison between actual and counterfactual outcomes to predict future decisions. For example, in the models of Selten et al. (2005), and Hart and Mas-Colell (2000), changes in the attraction of actions occur only when the outcome of the current action is inferior to the outcome that would have been obtained had an alternative action been taken, namely, when regret for foregone payoffs is likely (see also Grosskopf, Erev, & Yechiam, 2006; Yechiam & Busemeyer, 2006). In contrast, in the experience-weighted attraction model (Camerer & Ho, 1998, 1999; Ho et al., 2007) and in Selten and Chmura's (2008) model, all differences are assumed to have an effect. The models also differ in the weight assigned to actual versus counterfactual outcomes (Camerer & Ho, 1999) and in the weight assigned to positive versus negative outcomes (Selten & Chmura, 2008).

A second view of decision making holds that decisions are not a result of learning, but rather are constructed on the basis of an analysis of the situation. Such analysis may be based solely on the likelihood of different outcomes and the magnitude of rewards associated with them (Nash, 1951) or may also include the affect likely to arise with the outcomes of every possible option (e.g., Mellers, Schwartz, & Ritov, 1999; Ritov, 1996). An implication of this view is that once a situation is known and can be assumed not to change with time, decisions would be immune to the outcomes of recent decisions.

The present work belongs in the tradition of the former approach in that we assumed ongoing experience affects decisions. Specifically, it falls within the large and fast-growing body of research dealing with experience-based, repeated decision making (e.g., Barron & Erev, 2003; Busemeyer & Stout, 2002; Camerer & Ho, 1998, 1999; Erev & Barron, 2005; Hertwig, Barron, Weber, & Erev, 2004; Roth & Erev, 1995; Rakow & Newell, 2010; Selten et al., 2005;Yechiam & Busemeyer, 2006). However, our undertaking differs from most current ones in two important respects.

First, we were looking for the signature of impulsive reactions beyond any incremental change in behavior. We therefore focused on the effect of the outcome of the most recent decision on the upcoming decision. Recent models of repeated choices already argued that it is the reaction to only a few previous outcomes, rather than the whole learning history, that affects subsequent decisions (Erev & Barron, 2005; Erev et al., 2010; Selten & Chmura, 2008). Our approach is more extreme in that we focus on the exclusive effect of the most recent outcome. To the extent that outcomes induce impulsive reactions, we expected to find short-lived, persistent, effects of the most recent outcome, effects that would not diminish with experience.

Second, we used a novel approach in analyzing the data. Many current theories take a theoretical, top-down approach: They start by postulating psychologically motivated free parameters whose values are then estimated by fitting experimental data. The soundness of various psychological interpretations is assessed, mainly, by comparing the goodness of the fit of different models. In this approach, the number of possible psychological models can be very large (see, e.g., Hau, Pleskac, Kiefer, & Hertwig, 2008). In contrast, we adopted an empirical, bottom-up approach: We used the various outcome configurations of a given setup as the independent variables in predicting behavior. When the possible actions and rewards are both discrete, the possible outcomes of a situation are limited in number and well defined. In this approach, psychological interpretation is offered at a later stage, after effects have been demonstrated.

In our search for impulsive reactions we employed two very different tasks: an individual decision-making task and a two-person, strategic decision-making task. Both tasks called for repeated binary decisions; decisions were immediately rewarded, and information about the counterfactual reward—what would have been obtained had the alternative decision been made—was always provided. The first task was a variation on the popular two-armed-bandit game and involved a single player facing a random mechanism. The second task was the Parasite Game (Avrahami, Güth, & Kareev, 2005), which involved two competing players and a random mechanism. Two versions for each of the two tasks were employed.

In the first version of the individual decision-making task, a coin was hidden in one of two boxes. The player then selected one of the boxes, both boxes were opened, and the player received the coin if it was in the selected box. The coin was more likely to be in one of the boxes than the other—a fact that players could only learn from experience. Although normatively in this task the box that is more likely to contain the coin should always be selected (i.e., in game-theoretic terms it has a solution in a pure, maximizing strategy), previous research with both humans and other organisms has revealed that instead of maximizing, players make choices that roughly match the probabilities of reward for the two options (e.g., Estes, 1964, 1976; Herrnstein, 1990; Herrnstein & Vaughan, 1980; Vulkan, 2000). This version was employed to assess, in a minimal, standard task, the effect of the outcome of a decision on the subsequent choice and to explore changes that might occur with experience.

A drawback of this version is that the outcomes of actions taken are fully correlated with counterfactual outcomes: An outcome of no gain is also an outcome in which the alternative action would have resulted in gain and vice versa. To disentangle the effects of the actual and counterfactual outcomes, we designed a second version of the task in which the presence of a coin was independently determined for each box. Hence, a coin could be found in either box, in both boxes, or in neither box. Because a gain did not necessarily correspond to a better decision, and no gain did not necessarily correspond to a worse one, we could use this version to assess the effects of counterfactual outcomes.

The second task, the Parasite Game, involved two players who chose between two options and a stochastic mechanism that randomly drew between the same two options. A player’s reward was determined by the correspondence between the player’s own choice, the "move" of the stochastic random device, and the choice of the other player. A game-theoretic analysis shows that, unlike the first task, which normatively calls for a pure strategy, the optimal strategy in the Parasite Game is mixed.

The Parasite Game was introduced to detect the effect of impulsive reactions and to test the generality and utility of our approach in a different, richer, setting. The game also enabled us to compare the effects of impulses in the same task under different conditions: We used random pairing of players in one version and fixed pairing in the other. We could thus compare the impulses of players when they were and were not cognizant of their opponents’ outcomes, and hence, of their opponents' potential impulses.

THE INDIVIDUAL DECISION-MAKING TASK

In the individual task, a player repeatedly selected which of two boxes to open. After each decision, both boxes were opened, and their contents revealed. The player was rewarded if there was a coin in the box selected. One of the two boxes was more likely than the other to contain the coin. We call this task the Coins in Boxes Game. We chose to begin with a very simple setup that had only two possible outcomes—gain and no gain—to find out how well behavior can be predicted by the outcome of the most recent choice and whether the effects of the most recent outcome diminish as the player gains experience.

Experiment 1: The Coins in Boxes Game: Version I

Method

Materials and procedure. The Coins in Boxes Game was available on the Internet but limited to computers that were connected to the Hebrew University's servers. Before starting, players had to enter their Hebrew University student ID; the same ID could not be used to play the game more than once. The ID was later required to collect the payment earned. A time limit on each round ensured that players did not have time to take notes or make calculations and that the game was played in one sitting. The game was played for 100 rounds. Participants earned 1 point for each coin found; points were accumulated and converted to a monetary payment of 0.2 New Israeli Shekel (NIS) per point (approximately 5 U.S. cents at the time). The instructions stated explicitly that the probability with which a coin would be found in each box was set in advance and would remain unchanged throughout the experiment.

Design. On each round, a coin was found either in one box (with a probability of .60) or in the other (with a probability of .40). Thus, the presence of a coin in one box was perfectly correlated with the absence of a coin in the other box. The side (left or right) of the better box (i.e., the box more likely to contain the coin) and the actual sequence were both randomly determined, separately, for each player.

Participants. Forty-four students at the Hebrew University played this version of the game.

Results

All analyses were carried out on the decisions made from the eighth round onward. We did not include the first seven rounds because we intended to compare predictions based on the most recent outcome with predictions based on sequences of up to seven previous outcomes and had to base the performance of all models on the same data.

Overall, the better box was chosen with a probability of .65. This proportion was significantly greater than .50, $t(43) = 4.20$, $p < .001$, which indicates that choices were in line with the difference between the boxes. However, players were far from maximizing their reward: The probability of choosing the better box was also different from 1.00. Of the 44 participants, 6 chose the same box in all rounds. Of these, 4 always chose the better box, and 2 always chose the lesser one. To test for change with experience, we calculated the likelihood of choosing the better box separately for the first and last halves of the rounds. The proportions were .63 and .67, respectively. Although the proportion was higher in the second half than in the first half, choices in the second half were still far from the maximizing strategy. The number of players who chose the same box throughout was 8 in the first half and 9 in the second half; 2 players consistently chose the lesser box in both halves. Thus, even in the more advanced stage of the game, only a minority of the players adopted the maximizing strategy.

To find out to what extent the outcome of the most recent decision affected the upcoming decision, we conducted a linear regression in which the decision and outcome of a round were used to predict the subsequent decision, separately for each participant.[2] In this analysis, we constructed a dummy variable for each of the two possible outcomes—gain and no gain—incorporating the previous choice by assigning to the relevant dummy a value of +1 if the outcome occurred after a choice of the better box and a value of −1 if it occurred after a choice of the lesser box. These dummy variables were used to predict the choice of the better box. The analysis revealed that these two variables predicted the subsequent decision well (mean $R^2 = .26$, median $p < .001$). Thus, participants were differentially reactive to the different outcomes.

To see if the reaction to the outcomes made psychological sense, we looked at the predicted tendencies to choose the better box following each of the four contingencies. The predicted tendency to choose the better box was .82 following gain in the better box, .59 following no gain in that box, .36 following gain in the lesser box, and .58 following no gain in that box. Thus, the tendency to *repeat* a choice of the better box was .82 following gain and .59 following no gain, and the tendency to *repeat* a choice of the lesser box was .64 following gain and .42 following no gain. These tendencies, as expected, revealed that the outcome of the most

---

[2] A regression analysis cannot handle cases with a constant value for all predictors, and results in an undefined value in such cases. Hence, players who always chose the same box were excluded from this analysis.

recent round exerted an effect on the choice that followed: The tendency to repeat a choice was higher following gain than following no gain. This was true both for the better and for the lesser box. In addition, for each box, the tendency to repeat its choice was higher, on average, than the overall tendency to choose that box; in other words, inertia was apparently at play. Indeed, when considered alone, the most recent choice proved to have considerable predictive power (mean R^2 = .130, median p = .008).[3] Importantly, the results indicate a higher tendency to repeat choosing the better box—irrespective of outcome—than to repeat choosing the lesser box. This result demonstrates that participants realized that one box was better than the other, hence chose it more often; at the same time, the fact that the outcome of the most recent choice still affected their subsequent choice must be an indication of impulsive reaction.

To find out whether the effect of the most recent choice and its outcome diminished with experience, we conducted similar regressions separately for the two halves of the rounds. The effect of the previous outcome on the current choice not only did not diminish with experience, but in fact increased; R^2 rose from .21 for the first half to .31 for the second, $t(32) = -2.42$, p = .021. Given that the actual most recent choice was integrated into the predictors, this increase may have reflected an increase in inertia. Analyzing inertia separately for the two halves of the game revealed that although the effect of inertia rose too (R^2 = .09 for the first half and .12 for the second half), $t(32) = -1.36$, p = .183, this rise was insufficient to account for the overall rise in R^2. Thus, the effect of the previous outcome may indeed have increased between the halves.

Was any predictive power lost in focusing on the most recent outcome rather than a larger number of previous outcomes? To answer this question, we conducted a series of linear regression analyses that used a greater number of recent outcomes (from the two most recent to the seven most recent, as well as all outcomes previous to the current move). Three different memory-decay functions were employed for each of these analyses: no decay, power decay, and exponential decay. The goodness of prediction using larger numbers of previous outcomes was always lower than that obtained using the most recent outcome alone. The superiority of the model based on the most recent outcome may have resulted, however, from the fact that it directly represented the strong effect of inertia, whereas this effect was diluted in the models that relied on a longer sequence of outcomes. To compensate for this dilution, we conducted again all 21 regression-analyses, this time adding the actual previous choice as an additional,

[3] The finding of inertia is consistent with previous theories and models that acknowledge and incorporate the effect of inertia has been acknowledged and incorporated into corresponding models (e.g., Hart, 2005; Mookherjee & Sopher, 1994; for a review, see Camerer, 2003, chap. 6).

separate, predictor to all the models with more than one previous outcome considered. The addition of this parameter considerably improved the prediction of those models, but since adding this predictor increased the number of parameters by one, it had to be taken it into account. To do this, we used the Bayesian Information Criterion (the BIC measure, Raftery, 1995; Schwarz, 1978) and compared the BIC values based on the most recent outcome to each of the other 21 BIC values—separately for every player. The BICs based on the most recent outcome proved superior for a majority of the players in every one of the comparisons. We can thus conclude that no predictive power was lost in using the outcome of the most recent decision.

The results provide clear answers to our questions: The outcome of the single most recent decision predicted the subsequent decision well, and this effect did not diminish with experience.

This experiment revealed an effect of the outcome—gain or no gain—on the upcoming choice. As we explained earlier, however, the experiment could not reveal whether counterfactual outcomes had an effect on decisions, because reinforcement was fully correlated with the counterfactual outcomes: Gain was always associated with the satisfaction of having chosen well, and no gain was always associated with the regret for not having chosen otherwise. The second version of the Coins in Boxes Game addressed this issue. In that version, the presence of a coin was separately and independently determined for each box. This resulted in two new outcomes— a coin in both boxes or a coin in neither box—in addition to the outcomes that characterized Version I. Considerations of gain versus no gain only would still predict a higher tendency to repeat a choice that resulted in a gain than a choice that resulted in no gain, irrespective of what was in the other box; considerations of regret only (or of best-response only) would predict that neither of the two new outcomes would have an effect on the upcoming choice.

Experiment 2: The Coins in Boxes Game: Version II

In the second version of the Coins in Boxes Game, the probability of a coin being in the better box was again .60, and the probability of a coin being in the lesser box was again .40. Because in this version the presence of the coin was determined independently for each box, the probability that a coin would be found only in the better box was .36, and the probability that a coin would be found only in the worse box was .16 (we call these two cases the *either-or* cases). The probability of each of the two other possible outcomes—a coin in both boxes and a coin in neither box (the *both* and *neither* cases, respectively)—was .24. A different way to conceive of this version is as follows: There were two types of outcomes with gain)—one whose counterfactual outcome did differ and one whose counterfactual outcome did not (with the former possibly giving rise to satisfaction with one's choice). There were also two types of outcomes with no gain: again, one whose counterfactual outcome did differ and one whose counterfactual outcome did not (with the former possibly giving rise to regret). To study if and how the counterfactual outcomes affected the reactions to gains, we compared the tendencies of choice both within Version II and across the two versions of the game. We expected responses following the either-or cases to be similar to those in Version I, but considered the possibility that they would differ in the presence of the both or neither cases.

Method

The method for this experiment was identical to that for the experiment using Version I. Fifty-one students participated. The two experiments were conducted in parallel, with players randomly assigned to one of them.

Results

Players chose the better box with a probability of .67. As in Version I, this proportion was significantly different from .50, $t(50) = 5.06$, $p < .001$—in correspondence with the difference between the boxes—but far from 1.00. Of the 51 participants, 5 chose the same box in all rounds. Of these, 3 always chose the better box, but the other 2 always chose the lesser box. We again calculated the probabilities of choosing the better box separately for the first and second halves of the rounds; these probabilities were .61 and .72, respectively. Seven players chose the same box throughout the first half, and 11 chose the same box throughout the second half; 2 consistently chose the lesser box throughout both halves. Again, even in the more advanced stage of the game, only a minority of the players adopted the maximizing strategy.

A linear regression analysis with the current choice and outcome predicting the subsequent choice was conducted. It revealed, again, that these factors had considerable predictive power (mean $R^2 = .31$, median $p < .001$). In the cases comparable to the possibilities in Version I (i.e., rounds following either-or cases), the predicted tendency to choose the better box was .81 following gain in the better box, .60 following no gain in that box, .35 following gain in the lesser box, and .56 following no gain in that box. Note the similarity between these values and those in Version I (.82, .59, .36, and .58, respectively). More interesting are the corresponding values following rounds when the choice did not matter (i.e., following both and neither cases). On these rounds, the probability of choosing the better box was .79 following gain in the better box, .66 following no gain in that box, .38 following gain in the lesser box, and .50 following no gain in that box. These latter values are obviously less extreme: The effect of experiencing a gain or no gain was weaker when knowing that an alternative choice would not have changed the result. We subjected these value to a repeated measures analysis of variance with actual outcome (gain vs. no gain) and counterfactual outcome (different vs. same) as independent variables and found, in addition to a strong main effect of actual outcome, $F(1, 45) = 21.70$, $p < .001$, $\eta_p^2 = .325$, a significant interaction between actual and counterfactual outcome, $F(1, 45) = 4.14$, $p = .048$, $\eta_p^2 = .084$.

The choice probabilities following all contingencies in the two versions of the Coins in Boxes Game are presented in Figure 1. Clearly, although the effects of gain and no gain are evident under all contingencies, these effects were less pronounced when participants did not have the satisfaction of having chosen well or regret for having chosen badly.

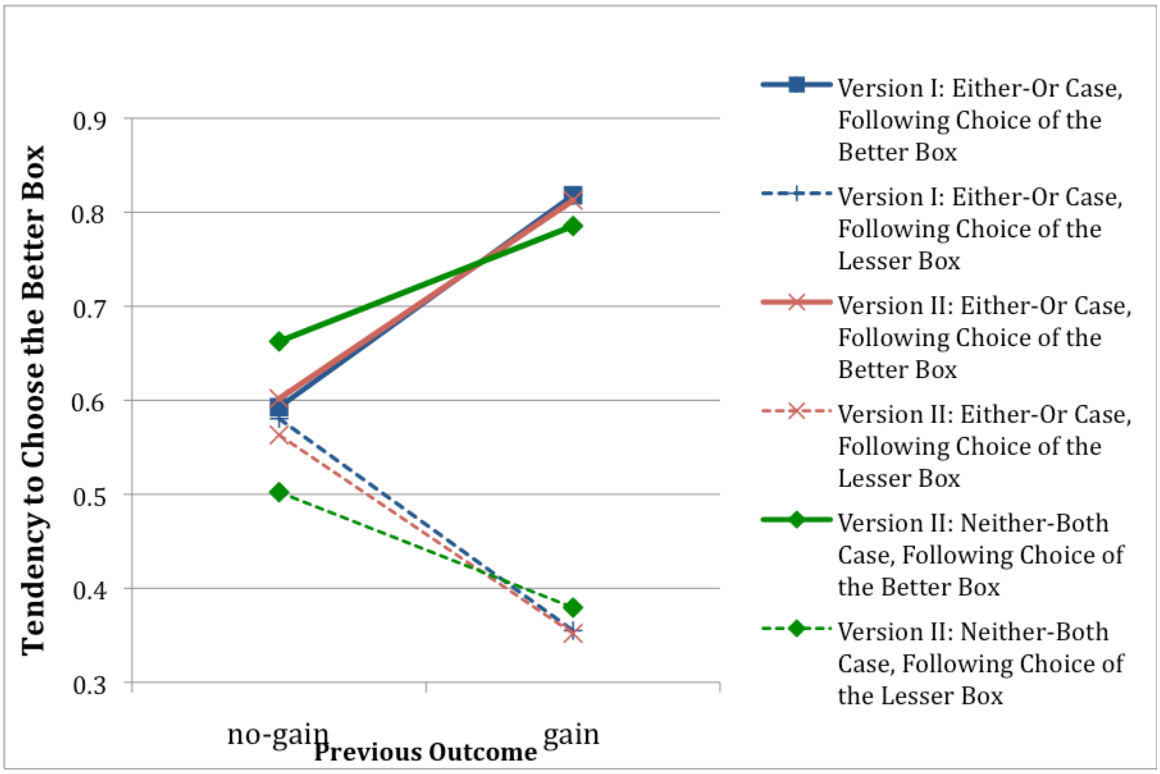

Figure 1. Tendency to choose the better box in the Coins in Boxes Game as a function of the previous outcome. Results are shown for decisions following either-or cases in Versions I and II and for decisions following both and neither cases in Version II. In the either-or cases, the coin was either in the chosen box (gain) or in the other box (no gain); in the 'both' cases, the coin was in both boxes (gain); and in the 'neither' cases, the coin was in neither box (no gain).

We again tested for inertia, predicting a choice on the basis of the preceding choice alone. As in Version I, inertia proved to have considerable predictive power (mean $R^2 = .18$, median $p = .002$).

Comparing predictability in the two halves revealed, again, that it did not diminish with experience (R^2 was .30 for each half). This was the case even though in this version inertia decreased in the second half (R^2s for inertia were .13 and .06 in the first and second half, respectively). Thus, as was the case for Version I, the effect of the most recent outcome not only did not diminish with experience, but may have even increased slightly.

As before, we tested if any predictive power was lost in focusing on the most recent outcome rather than a larger number of previous outcomes. We again used outcome sequences ranging from two to seven previous rounds, as well as the entire set of rounds preceding the current move, and employed the same three memory decay functions (no decay, power, and exponential decay). The results showed that predictions based on the most recent outcome again outperformed every other combination. When adding the most recent choice as an extra parameter to the models with more than one previous outcome, calculating the BIC values, and comparing those values for every player, we found that the BICs based on the most recent outcome proved superior for a majority of the players in every comparison. We therefore conclude that here too no predictive power was lost in using the outcome of the most recent decision.

The results of Version II replicated those of Version I in showing that the outcome of the most recent choice has a significant effect on the following choice, an effect that does not diminish with experience. The results of Version II extend those of Version I in showing how various counterfactual outcomes interact with actual ones—attenuating or amplifying the effects of actual outcomes depending on the correspondence between them.

So far, we have reported on choice behavior in a setting in which a single player faced a stochastic mechanism. We now turn to the Parasite Game, a more complex setup in which payoffs were determined not only by the player's own decision and the draw of a stochastic mechanism, but also by the decision of a competitor.

THE STRATEGIC DECISION-MAKING TASK

The Parasite Game models situations in which the successful choice of only one player is necessary and sufficient to produce a reward, but the beneficiary of the reward is determined by the choice of the other player. For example, consider the fish in a river with two tributaries. Because of changes in river currents, these fish are to be found in only one of the river's tributaries on a given day. On every day, a fisherman chooses the tributary where he will fish, and a robber chooses the tributary where he will prowl. Fish are caught only if the fisherman chooses the tributary where the fish are to be found that day. The fisherman gets fish only if he fishes at the right tributary and the robber is not there; the robber gets fish only if he is present at the same tributary as the fisherman and the fish.

The Parasite Game involves two players and a stochastic mechanism. Each player chooses privately between two options; simultaneously, the stochastic mechanism randomly draws between those options, which have probabilities w and $1 - w$ ($w \geq .5$). Payoffs are determined as follows: Player A wins if his or her choice corresponds to the draw of the stochastic mechanism and that of player B does not, player B wins if his or her own choice corresponds both to player A's choice and to the draw, and no one wins if player A misses the draw (see Figure 2, which presents the game tree and payoff matrix).

		Player B	
		H	T
Player A	H	0, w	w, 0
	T	1-w, 0	0, 1-w

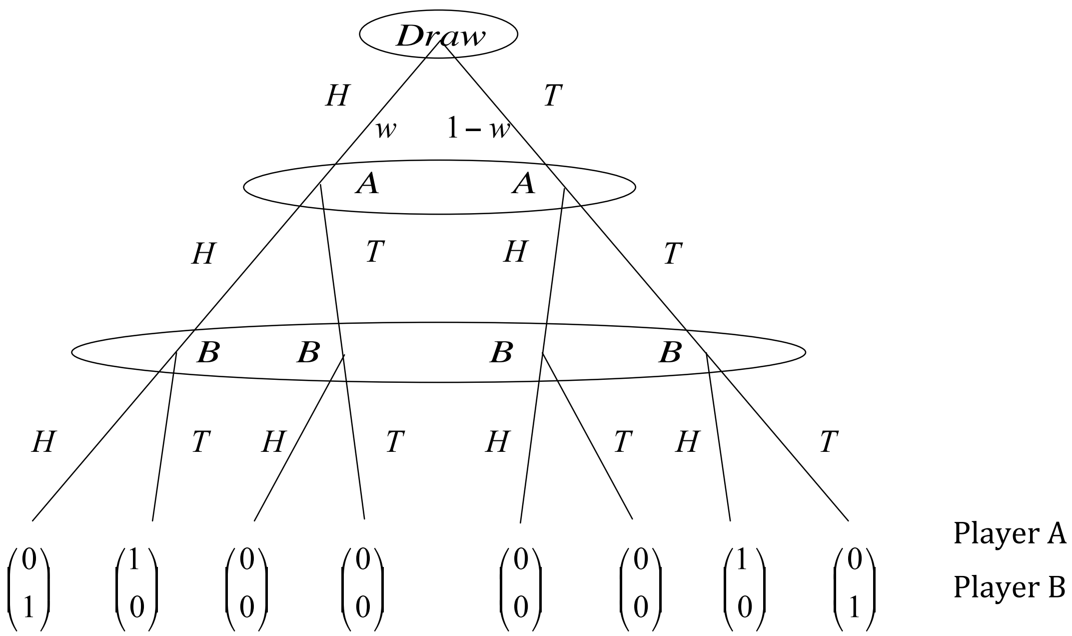

Figure 2. Payoff matrix and game tree of the Parasite Game. "*w*" is the probability with which the stochastic mechanism draws option *H*. With $w > .5$, option *H* is the one more likely to be drawn, and option *T* the less likely. The numbers at the end nodes represent the payoffs for players A and B

Given that the stochastic mechanism is more likely to draw one of the options than the other, the decision of which option to choose becomes quite complex for both players: It is clear that always choosing the same option, namely, adopting a pure strategy, would be to either player's disadvantage; a mixed strategy is called for, but the makeup of the mixture is not easy to figure out.[4] As mentioned, player A wants to match the draw of the stochastic mechanism and avoid player B's choice, and player B wants to match both the draw of the stochastic mechanism and player A's choice. Thus, for each player, a choice can turn out to be good or bad vis-à-vis the draw of the stochastic mechanism (matching is good for both players) and good or bad vis-à-vis the opponent's choice (with "good" meaning a nonmatch for player A and a match for player B). A player wins only in the case of one of these four outcomes—when his or her choice turns out to be good vis-à-vis both the draw of the stochastic mechanism and the other player's choice. As for the three other outcomes, though all result in no gain, they differ with respect to their counterfactual gain. In one case, when the choice is bad both vis-à-vis the draw and vis-à-vis the opponent's choice, having chosen the other alternative would have resulted in a gain. In this case, the player may regret not having chosen otherwise. In the remaining two cases, choosing differently would only have moved the player from one of these cases to the other case; because neither would result in a gain, there is no reason for regret.

We used players' choice and the outcome on the most recent round to predict players' subsequent choice. Two versions of the game were employed: In the first, opponents were randomly paired on each round (the *stranger* design), and in the second, players were paired with the same opponent throughout the game (the *partner* design).

[4]There are two mixed-strategy game-theoretic solutions for this game: a unique equilibrium (Nash, 1951) and an impulse-balance equilibrium (Selten et al., 2005). The latter equilibrium is the resting point of the Learning Direction Theory (Selten & Stoecker, 1986; Selten et al., 2005). The theory postulates that when a comparison of the outcome of a decision and the counterfactual outcome reveals that the latter would have been better than the former, an impulse away from that decision takes place. Behavior will end up at a point in which the impulses from all options are balanced, thus reaching equilibrium. As we have shown elsewhere (Avrahami et al., 2005), if the stochastic mechanism draws Option H with probability w and Option T with probability $1 - w$, the Nash equilibrium is for player A to choose Option H with probability $1 - w$ and for player B to choose that option with probability w. The impulse-balance equilibrium is for player A to choose Option H with probability .5—irrespective of w—and for player B to choose that option with probability w.

Experiment 3: The Parasite Game: Version I

Method

Materials, procedure and design. The game was set up as a computerized "color game" in which the computer randomly drew either red or green, and each player chose between red and green. The draw and both choices were then revealed, and the winner, if there was one, was determined. A win was worth 1 NIS.

In each of two sessions, 12 players played the game for 100 rounds. Role was manipulated between participants. In every session, 6 players, chosen at random, were assigned to the role of player A, and the remaining 6 players were assigned to the role of player B. Roles were fixed for the duration of the game. One of the colors was drawn with a probability of .75, and the other with a probability of .25. The more frequent color—red or green—was counterbalanced across sessions. Players were randomly paired on each round.

Players were told which combinations of draws and choices led to gains and no gains, and that they would play against an anonymous opponent who was randomly determined on each round. The probability with which each color was drawn by the stochastic mechanism was not announced. The draw was the same for all players in a session. The game was conducted in a computer lab (The Ratiolab) located at the psychology department of the Hebrew University of Jerusalem. Each player was seated in a separate cubicle facing a personal computer.

Participants. Twenty-four students at the Hebrew University participated for a monetary reward.

Results

We first calculated the probability with which each player chose the more frequent color. Participants in the role of player A chose the more frequent color with a probability of .51, and participants in the role of player B chose that color with a probability of .75.[5] Analyzing the two halves of the game separately revealed that participants in the role of player A chose the more frequent color with probabilities of .59 and .45 in the first and second halves, respectively, and those in the role of player B chose that color with probabilities of .68 and .81 in the first and second halves.

[5] With the computer drawing the more frequent color with a probability of .75, the observed numbers correspond very closely to the .50 and .75 predicted by the impulse-balance equilibrium (Selten, et al. 2005) see footnote 4 above.

We examined effects of the outcome of the most recent decision on the subsequent one by regressing players' decisions on the choice and outcome in the most recent round. Unlike in the individual decision task, in which only the outcome of the stochastic mechanism and the choice of the player were considered, the choice of the opponent was also part of the contingency of every round in the Parasite Game. Thus, there were four outcomes. The mean R^2 of the regression was .23, median $p < .001$, indicating that choice could be predicted well on the basis of aspects of the most recent round.

To see if reactions to the outcome of the most recent round were in line with what would be expected from a psychological interpretation, we looked at the tendency to choose the more frequent color following each of the contingencies. The tendency to choose the more frequent color immediately after having chosen that color was highest following gain (good choice vis-à-vis both the draw and the opponent's choice, .80). The tendency to choose the more frequent color again was lowest following a choice that could have resulted in a gain had the alternative color been chosen (bad choice vis-à-vis both the draw and the other player's choice, .56), and intermediate following choices for which the alternative choice would have made no difference (.72 for choices good vis-à-vis the draw and bad vis-à-vis the opponent's choice, .75 for choices bad vis-à-vis the draw and good vis-à-vis the opponent's choice). The complementary pattern was evident, of course, after participants chose the less frequent color; the comparable values for choosing the more frequent color were .37, .60, .45, and .42. As was the case in the Coins in Boxes game, the overall greater tendency to repeat the more frequent color than the other, may reflect strategic considerations (by player B in particular). Still, we believe that the difference in those tendencies following different outcomes reflects impulsive reactions to those outcomes.

These patterns of choice were clearly in line with the outcomes. A repeated measures analysis of variance on the predicted probability of choosing the more frequent color, with outcome of the previous round as a within-players factor and role as a between-players factor, revealed a significant effect of outcome, $F(3, 66) = 13.15$, $p < .001$, $\eta_p^2 = .374$. The effect of outcome did not interact with role, $F < 1$. The overall pattern of results clearly exhibited both an effect of gain and an effect of the counterfactual result: The tendency to repeat a choice was higher following a gain than following no gain and was lowest following a choice that may have induced regret for not having made the other choice. This pattern was captured by a highly significant linear contrast $F(1, 22) = 25.10$, $p < .001$, $\eta_p^2 = .533$.

As in the individual decision task, players exhibited inertia, the tendency to repeat the previous choice irrespective of outcome. The probability of choosing the more frequent color, averaged across outcomes, was higher following a choice of that color than following a choice of the infrequent color. Indeed, choice on a given round was predicted well by the most recent choice alone mean $R^2 = .15$, median $p = .007$. The tendencies to choose the more frequent color following each of the outcomes are presented in Figure 3.

Did the effect of outcome diminish as players gained experience? Separate regressions for the first and second halves of the game revealed that the magnitude of R^2 diminished slightly, from .25 to .21, but not significantly, $t(22) = 0.71$, $p = .48$.

As in the previous task, we compared the goodness of prediction based solely on the most recent outcome with the goodness of prediction based on longer sequences of previous outcomes, using various decay functions. Again, calculating the BIC values, and comparing those values for every player, we found that the BICs based on the most recent outcome proved superior for a majority of the players in all comparisons but one (in which the two models were equal). Once again, no predictive power was lost in using the outcome of the most recent decision.

Experiment 4: The Parasite Game: Version II

Method

Version II of the Parasite Game was identical to Version I with the exception of how opponents were paired. In Version II, players were paired once, at the beginning of the game, and played with the same, albeit anonymous, opponent for all 100 rounds. As a result, players knew what their opponents had experienced in the most recent round. Thus, considerations of the opponent's potential impulses, as well as awareness that the opponent could know one's own impulses, may have influenced choices. However, if participants differed in the depth at which they consider their opponents' likely reactions to impulses (exhibiting different *k* levels, Camerer, Ho, & Chong, 2004), such variety could have attenuated the average effects of outcomes on subsequent decisions.

Results

In this version, player A chose the more frequent color with a probability of .57, and player B chose that color with a probability of .67. Although these numbers were no longer in perfect correspondence with the impulse-balance equilibrium, they were further from the unique Nash equilibrium. The probability of choosing the more frequent color was .60 in the first half and .54 in the second half for player A, and .63 in the first half and .71 in the second half for player B.

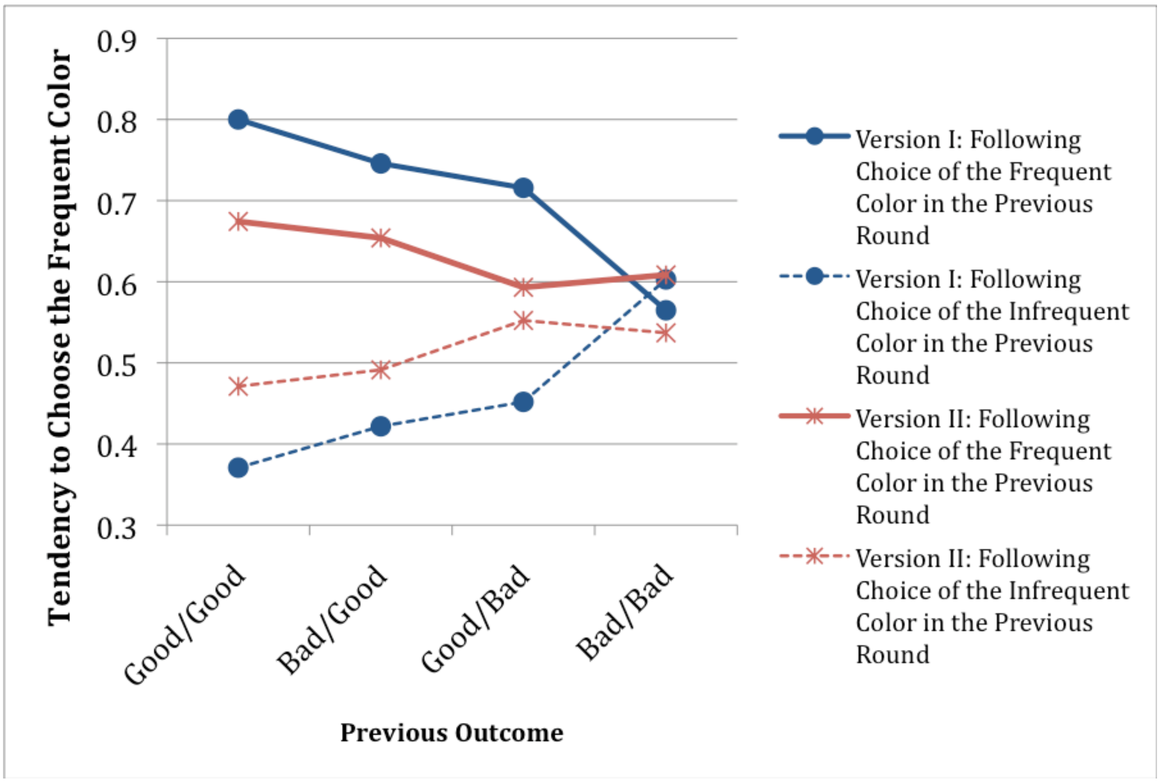

Figure 3. Predicted tendency to choose the better box in the Parasite Game as a function of the previous outcome. The labels along the x-axis indicate whether the previous choice was good or bad vis-à-vis the draw of the stochastic mechanism and vis-à-vis the choice of the opponent. For example, "bad/good" indicates a choice that was bad vis-à-vis the draw of the stochastic mechanism and good vis-à-vis the opponent's choice.

An important question that this version could answer was whether players made an attempt to be less predictable when they were always paired with the same opponent. The regression analysis that predicted the upcoming choice from the most recent choice and outcome revealed that players were still quite predictable (mean R^2 = .19, median p = .002), although predictability was somewhat lower than in Version I. There was some decrease in predictability from the first to the second half; mean R^2 declined from .26 to .19 (median ps = .029 and .183, respectively), but the difference did not reach significance, $t(21) = 1.36$, p = .189. Nevertheless, it seems that with practice, players may have started to suppress the effects of their impulses.

Interestingly, the average tendency to choose a color no longer corresponded to the implications of the outcome on the previous round. This may have resulted from players being aware of their opponents' ability to predict their impulsive reactions. The average tendency after having chosen the frequent color were: .67 following gain (good choice vis-à-vis both the draw and the opponent's choice), .61 following a choice that could have resulted in a gain had the alternative color been chosen (bad choice vis-à-vis both the draw and the other player's choice), and .59 and .65 following choices for which the alternative choice would have made no difference (for choices good vis-à-vis the draw and bad vis-à-vis the opponent's choice and for choices bad vis-à-vis the draw and good vis-à-vis the opponent's choice, respectively). The corresponding values after having chosen the infrequent color were .47, .54, .55, and .49. These values no longer differed significantly, $F(3, 63) = 1.20$, $p = .32$; in addition, the linear contrast was not significant, $F(1, 21) = 1.56$, $p = .23$. This finding, together with the relatively high success of the regression in predicting players' choices, indicates that although players were relatively predictable, different players reacted differently to the most recent outcome, with more or less regard to their opponents. The tendencies to choose the more frequent color following each outcome are presented in Figure 3.

Inertia was again evident: Predicting a choice on the basis of the previous one resulted in a mean R^2 value of .10, median $p = .102$. Though inertia was somewhat weaker in Version II than in Version I, the difference between versions did not reach significance $t(45) = 0.95$.

We again compared the goodness of prediction based on the most recent outcome with the goodness of prediction based on longer sequences of previous outcomes, using the three decay functions. Calculating the BIC values, and comparing those values for every player, we found that the BICs based on the most recent outcome proved superior for a majority of the players in every comparison. Once again, no predictive power was lost in using the outcome of the most recent decision.

GENERAL DISCUSSION

Ever since the early days of psychology, it has been known that the valence of the outcome of an action affects subsequent behavior: Actions with positive outcomes tend to be repeated, whereas actions with negative outcomes tend to be avoided. This fact, captured by Thorndike's (1927) law of effect, has been demonstrated in innumerable studies. Postulating the presence of short-lived impulsive reactions, we explored how the valence of the immediately preceding outcome affects behavior. Two issues were considered: The first involved the strength of the effects of impulses on upcoming decisions; the second involved the stability of the effect of impulses over time. In addition, we compared the explanatory power of the most recent outcome with that of the accumulation of a larger set of outcomes ranging further into the past. Our data, collected in both individual and strategic decision-making tasks, provide clear insights into these issues.

With respect to the first issue, our data indicate that by taking into account the last action taken and its outcome, it is possible to predict the following action quite well. The outcome of the previous action—positive or negative—and the outcomes of alternatives not taken—whether better or worse—had a strong effect on the subsequent action. We also found that part of the predictive power of the most recent action stems from a strong effect of inertia—a tendency to repeat a previous action. With respect to the second issue, our findings indicate that the effects of impulses do not diminish with experience. The effect of the immediate past on the following decision proved to be as strong in the later stages of the games as in the earlier ones. Finally, our data clearly indicate that predictions based on the most recent action were as good as or superior to predictions based on a longer past. This was true for various numbers of past moves (from two up to seven moves, or moves from all preceding rounds) and for a variety of decay functions modeling the effects of the past (no decay, power decay, or exponential decay). This finding complements that of a nondiminishing effect of impulses.

It is important to note that a nondiminishing effect of impulses does not imply that aggregate behavior would not change with experience. In fact, as a result of consistent and stable reactance to just the previous outcome aggregate behavior at both the group and the individual level would exhibit what are regarded as the two signatures of learning: a decelerating increase in the probability of choosing the better alternative and a stabilization of that probability at a level at which rewards are not necessarily maximized.

This is not to say that when intelligent and motivated people make consequential decisions, they do not go into complicated considerations before acting. For example, before investing money, people may read economics papers, search the Internet, discuss the options with friends, or seek professional advice. Their overall behavior would reflect the integration of all these sources of information. But even with all the reasoning and all the advice received, they would still, we argue, tend to react impulsively to recent market outcomes, often buying and selling prematurely. They may be aware of the rational considerations they employ in making decisions, but they are unaware of the effects of impulses on their behavior. When discussing the two games with students who experienced them, we were given a plethora of explanations of how they had played: a variety of reasons for playing one way or another, and even suggestions as to the cognitive procedure that had been followed. None of the students mentioned reactions to the outcome of the most recent decision. In fact, they were quite amazed to see the mark of such reactions when their own data were presented to them.

Finally, by comparing the impulses arising in a decision task under different conditions, one can reveal the affective reactions to different situations. We have demonstrated how such comparisons can be used for setups that differ in counterfactual outcomes or in the information available to players about their competitors. Analysis of the impulses arising in a task could also be used to compare players coming from different groups with well-defined characteristics (e.g., from healthy and unhealthy populations). The analysis of impulsive reactions to recent outcomes thus provides a new angle from which repeated decisions can be studied.

REFERENCES

Avrahami, J., Güth, W., & Kareev, Y. (2005). Games of competition in a stochastic environment. *Theory and Decision, 59,* 255–294.

Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making, 16,* 215–233.

Busemeyer, J., & Stout, J. C. (2002). A contribution of cognitive decision models to clinical assessment: Decomposing performance on the Bechara gambling task. *Psychological Assessment, 14*, 253–262.

Camerer, C. F. (2003). *Behavioral game theory: Experiments in strategic interaction.* New York: Russell Sage Foundation.

Camerer, C. F., & Ho, T. H. (1998). Experience-weighted attraction learning in coordination games: Probability rules, heterogeneity, and time-variation. *Journal of Mathematical Psychology, 42,* 305–326.

Camerer, C. F., & Ho, T. H. (1999). Experience-weighted attraction learning in normal form games. *Econometrica, 67,* 827–874.

Camerer, C. F., Ho, T. H., & Chong, J. K. (2004). A cognitive hierarchy model of thinking in games. *Quarterly Journal of Economics, 119,* 861–898.

Cournot, A. (1960). *Researches in the mathematical principles of the theory of wealth* (N. Bacon, Trans.). London: Haffner. (Original work published 1838)

Erev, I., & Barron, G. (2005). On adaptation, maximization, and reinforcement learning among cognitive strategies. *Psychological Review, 112*, 912–931.

Erev, I., Ert, E., Roth, A. E., Haruvy, E., Herzog, S. M., Hau, R., Hertwig, R., Stewart, T., West, R., & Lebiere, C. (2010). A choice prediction competition: Choices from experience and from description. *Journal of Behavioral Decision Making, 23,* 15-47.

Erev, I., & Roth, A. E. (1998). Modeling how people play games: Reinforcement learning in experimental games with unique mixed equilibria. *American Economic Review, 88,* 848–881.

Estes, W. K. (1964). Probability learning. In A. W. Melton (Ed.), *Categories of human learning* (pp. 89–128). New York: Academic Press.

Estes, W. K. (1976). The cognitive side of probability learning. *Psychological Review, 83,* 37–64.

Grosskopf, B., Erev, I., & Yechiam, E. (2006). Foregone with the wind: Indirect payoff information and its implications for choice. *International Journal of Game Theory, 34,* 285-302.

Hau, R. Pleskac, T. J., Kiefer, J., & Hertwig, R. (2008). The description-experience gap in risky choice: The role of sample size and experienced probabilities. *Journal of Behavioral Decision Making, 21*, 493–518.

Hart, S. (2005). Adaptive heuristics. *Econometrica, 73,* 1401–1430.

Hart, S., & Mas-Colell, A. (2000). A simple adaptive procedure leading to correlated equilibrium. *Econometrica, 68,* 1127–1150.

Herrnstein, R. J. (1990). Rational choice theory: Necessary but not sufficient. *American Psychologist, 45,* 356–367.

Herrnstein, R. J., & Vaughan, W., Jr. (1980). Melioration and behavioral allocation. In J. E. R. Staddon (Ed.), *Limits to action: The allocation of individual behavior* (pp. 143–176). New York: Academic Press.

Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science, 15,* 534–539.

Ho, T. H., Camerer, C. F., & Chong, J. K. (2007). Self-tuning experience-weighted attraction learning in games. *Journal of Economic Theory, 133,* 177–198.

Levine, M. (1966). Hypothesis behavior by humans during discrimination learning. *Journal of Experimental Psychology, 71,* 331–338.

Matsen, F. A., & Nowak, M. A. (2004). Win-stay, lose-shift in language learning from peers. *Proceedings of the National Academy of Sciences, USA, 101,* 18053–18057.

Mellers, B. A., Schwartz, A., & Ritov, I. (1999). Emotion-based choice. *Journal of Experimental Psychology: General, 128,* 332–345.

Milinski, M. (1993). Cooperation wins and stays. *Nature, 364,* 12–13.

Mookherjee, D., & Sopher, B. (1994). Learning behavior in an experimental matching pennies game. *Games and Economic Behavior, 7,* 62–91.

Nash, J. F. (1951). Non-cooperative games. *Annals of Mathematics, 45,* 286–295.

Nowak, M. A., & Sigmund, K. (1993). A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner's Dilemma game. *Nature, 364,* 56–58.

Raftery, A. E. (1995). Bayesian model selection in social research. *Sociological Methodology, 25*, 111-163.

Rakow, T., & Newell, B. R. (2010). Degrees of uncertainty: An overview and framework for future research on experienced-based choice. *Journal of Behavioral Decision Making, 23,* 1-14.

Ritov, I. (1996). Probability of regret: Anticipation of uncertainty resolution in choice. *Organizational Behavior and Human Decision Processes, 66,* 228–236.

Roth, A. E., & Erev, I. (1995). Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. *Games and Economic Behavior, 8,* 164–212.

Schwarz, G. (1978). Estimating the dimension model. *The Annals of Statistics, 6*, 461–464.

Selten, R., Abbink, K., & Cox. R. (2005). Learning direction theory and the winner's curse. *Experimental Economics, 8,* 5–20.

Selten, R., & Chmura, T. (2008). Stationary concepts for experimental 2x2-games. *American Economic Review, 98,* 938–966.

Selten, R., & Stoecker, R. (1986). End behavior in sequences of finite Prisoner's Dilemma supergames. *Journal of Economic Behavior and Organization, 7*, 47-70.

Thorndike, E. L. (1927). The law of effect. *The American Journal of Psychology, 39,* 212–222.

Vulkan, N. (2000). An economist's perspective on probability matching. *Journal of Economic Surveys, 14,* 101–108.

Yechiam, E., & Busemeyer, J. R. (2006). The effect of foregone payoffs on underweighting small probability events. *Journal of Behavioral Decision Making, 19,* 1–16.